

07-19-00

A

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU9-2000-0112-US1

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
 Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

A PLURALITY OF FILE SYSTEMS USING WEIGHTED ALLOCATION TO ALLOCATE SPACE ON ONE OR MORE STORAGE DEVICES

and invented by:

Wayne A. Sawdon, Roger L. Haskin, Frank B. Schmuck, James C. Wyllie

JC675 U.S. PRO
09/618508

07/18/00

If a CONTINUATION APPLICATION, check appropriate box and supply the requisite information:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Which is a:

Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. Filing fee as calculated and transmitted as described below
2. Specification having 55 pages and including the following:
 - a. Descriptive Title of the Invention
 - b. Cross References to Related Applications (*if applicable*)
 - c. Statement Regarding Federally-sponsored Research/Development (*if applicable*)
 - d. Reference to Microfiche Appendix (*if applicable*)
 - e. Background of the Invention
 - f. Brief Summary of the Invention
 - g. Brief Description of the Drawings (*if drawings filed*)
 - h. Detailed Description
 - i. Claim(s) as Classified Below
 - j. Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU9-2000-0112-US1

Total Pages in this Submission

Application Elements (Continued)

3. Drawing(s) (*when necessary as prescribed by 35 USC 113*)
a. Formal Number of Sheets Seven (7)
b. Informal Number of Sheets _____

4. Oath or Declaration
a. Newly executed (*original or copy*) Unexecuted
b. Copy from a prior application (37 CFR 1.63(d)) (*for continuation/divisional application only*)
c. With Power of Attorney Without Power of Attorney
d. **DELETION OF INVENTOR(S)**
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).

5. Incorporation By Reference (*usable if Box 4b is checked*)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied
under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.

6. Computer Program in Microfiche (*Appendix*)

7. Nucleotide and/or Amino Acid Sequence Submission (*if applicable, all must be included*)
a. Paper Copy
b. Computer Readable Copy (*identical to computer copy*)
c. Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. Assignment Papers (*cover sheet & document(s)*)

9. 37 CFR 3.73(B) Statement (*when there is an assignee*)

10. English Translation Document (*if applicable*)

11. Information Disclosure Statement/PTO-1449 Copies of IDS Citations

12. Preliminary Amendment

13. Acknowledgment postcard

14. Certificate of Mailing

First Class Express Mail (*Specify Label No.*): EL601699735US

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU9-2000-0112-US1

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. Certified Copy of Priority Document(s) (*if foreign priority is claimed*)

16. Additional Enclosures (*please identify below:*)

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	71	- 20 =	51	x \$18.00	\$918.00
Indep. Claims	7	- 3 =	4	x \$78.00	\$312.00
Multiple Dependent Claims (check if applicable)	<input type="checkbox"/>				\$0.00
				BASIC FEE	\$690.00
OTHER FEE (specify purpose)					\$0.00
				TOTAL FILING FEE	\$1,920.00

A check in the amount of to cover the filing fee is enclosed.

The Commissioner is hereby authorized to charge and credit Deposit Account No. 09-0463 (IBM) as described below. A duplicate copy of this sheet is enclosed.

- Charge the amount of \$1,920.00 as filing fee.
- Credit any overpayment.
- Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
- Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).

Blanche E. Schiller
Signature

Blanche E. Schiller, Esq.
 Reg. No. 35,670
 HESLIN & ROTHENBERG, P.C.
 5 Columbia Circle
 Albany, NY 12203
 Telephone (518) 452-5600
 Facsimile (518) 452-5579

Dated: July 18, 2000

CC:

CERTIFICATE OF MAILING BY "EXPRESS MAIL"

In Re Application of: Sawdon et al.

Title: A PLURALITY OF FILE SYSTEMS USING WEIGHTED ALLOCATION
TO ALLOCATE SPACE ON ONE OR MORE STORAGE DEVICES

Attorney Docket No.: POU9-2000-0112-US1

"EXPRESS MAIL" MAILING LABEL NO. EL601699735US

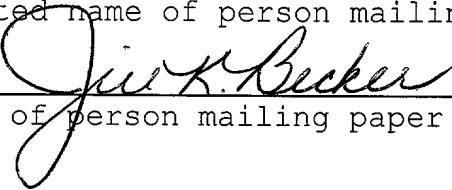
Date of Deposit July 18, 2000

I hereby certify that this paper is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and addressed to:

BOX PATENT APPLICATION
ASSISTANT COMMISSIONER FOR PATENTS
WASHINGTON, D.C. 20231

JILL K. BECKER

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

Enclosures:

- * Utility Patent Application Transmittal Letter (3 pages)
(in duplicate)
- * U.S. Patent Application which includes:
Specification (34 pages), 71 Claims (20 pages),
Abstract (1 page)
- * Seven (7) sheets of Formal Drawings
- * Declaration and Power of Attorney (unsigned) (4 pages)
- * Information Disclosure Citation (2 pages) and twelve (12)
references
- * Two (2) Acknowledgment Postcards

**A PLURALITY OF FILE SYSTEMS USING
WEIGHTED ALLOCATION TO ALLOCATE SPACE
ON ONE OR MORE STORAGE DEVICES**

Cross-Reference to Related Applications

5 This application contains subject matter which is related to the subject matter of the following application/issued patent, each of which is assigned to the same assignee as this application. Each of the below listed applications/patents is hereby incorporated herein by reference in its entirety:

10

“Determining The Order And Frequency In Which Space Is Allocated On Individual Storage Devices,” Sawdon et al., (Docket No. POU9-2000-0111-US1), Serial No. _____, filed herewith; and

15 “Parallel File System And Method With Allocation Map,” Schmuck et al., U.S. Patent No. 5,960,446, Issued September 28, 1999.

Technical Field

20 This invention relates, in general, to allocating space on storage devices, and in particular, to enabling a plurality of file systems to use weighted allocation to allocate space on one or more storage devices.

Background Art

Many computing environments include file systems, which enable other application programs to store data on and retrieve data from storage devices. In particular, a file system allows application programs to create files and to give them names (a file is a named data object of arbitrary size), to store (or write) data into files, to read data from files, to delete files, and to perform other operations on files.

A file structure is the organization of data on the storage devices. In addition to the file data itself, the file structure contains meta data, which includes, for instance, the following: a directory that maps file names to the corresponding files; file meta data that contains information about the file, including the location of the file data on the storage device (i.e., which device blocks hold the file data); an allocation map that records which device blocks are currently in use to store meta data and file data; and a superblock that includes overall information about the file structure (e.g., the locations of the directory, allocation map, and other meta data structures).

In order to store successive data blocks of a file to distinct devices, such as disks or other storage devices, a technique known as striping is used. Striping may also be used to store the file system's meta data. The advantages of striping include high performance and load balancing. In

striping, the file system writes successive blocks of a
file, or the file's meta data, to distinct devices in a
defined order. For example, the file system may use a
round-robin allocation, in which successive blocks are
5 placed according to a cyclic permutation of the devices.
This permutation is called the stripe order. The stripe
order defines the order and frequency of allocations (and
thus, writes) to each device in the file system. For
example, a system with four disks using a simple round-robin
10 allocation scheme would allocate space on each disk in
consecutive order, namely: 1, 2, 3, 4, 1, 2, 3, 4....

This simple round-robin allocation is used by most
striped file systems for allocation. Although, round-robin
allocations may be sufficient in some circumstances for a
15 system that includes homogeneous devices, it proves to be
inadequate for a system with heterogeneous devices, and it
proves to be inadequate for various circumstances in which
homogeneous devices are used.

As one example, a round-robin allocation is inadequate
20 for devices of different storage capacities or throughput.
Under round-robin allocation, all devices are allocated
equally. Consequently, subsequent access to the data is
typically spread equally across the devices as well. For
systems that include devices with different storage
25 capacities, the small devices fill before the larger devices
and then, must be excluded from the stripe order, thus
reducing the parallelism and performance for all subsequent
writes. Furthermore, the data striped across the reduced

set of devices has reduced performance for all subsequent accesses.

Likewise, for systems that include devices with different throughput, round-robin allocation fails to maximize the throughput for allocation and all subsequent accesses to the data. Additionally, round-robin allocation has no capability for rebalancing a system that is in an unbalanced state. An unbalanced state can occur for a variety of reasons including, for instance, when devices are partitioned between files or operating systems; when empty devices are added to an existing file system; or when the allocation policy changes. To rebalance such a system, extraordinary measures are required by the user, such as restriping of all the data in the file system.

Striping can be performed by a single file system, or by a plurality of file systems of a shared device file environment (e.g., a parallel environment). In a shared device file environment, a file structure residing on one or more storage devices is accessed by multiple file systems running on multiple computing nodes. A shared device file environment allows an application (or job) that uses the file structure to be broken up into multiple pieces that can be run in parallel on multiple nodes. This allows the processing power of these multiple nodes to be brought to bear against the application.

The above-described problems associated with striping are exacerbated in a parallel environment. Thus, a need

still exists for a parallel allocation technique that is general enough to be used in a wide variety of circumstances. Further, a need exists for a capability that enables rebalancing of the allocations to better match the
5 current conditions and requirements of the system and/or devices.

Summary of the Invention

The shortcomings of the prior art are overcome and additional advantages are provided through the provision of
10 a method of managing the allocation of space on storage devices of a computing environment. The method includes, for instance, obtaining one or more weights for one or more storage devices of the computing environment; and allocating space on at least one storage device of the one or more
15 storage devices in proportion to at least one weight obtained for the at least one storage device, wherein the allocating is performed by a plurality of file systems of the computing environment.

In a further embodiment, a method of managing the allocation of space on storage devices of a computing environment is provided. The method includes, for instance, obtaining a weight for each storage device of at least a subset of storage devices of a plurality of storage devices of the computing environment; and allocating space on each
20 storage device of the at least a subset of storage devices in proportion to the weight assigned to the storage device,
25 wherein the allocating is performed by a plurality of file

systems such that each file system of the plurality of file systems allocates space on one or more storage devices of the at least the subset of storage devices.

5 System and computer program products corresponding to the above-summarized methods are also described and claimed herein.

The capabilities of one or more aspects of the present invention advantageously provide for the allocation of space, by a plurality of file systems, across one or more 10 storage devices, such that the space on each device is allocated and thus, consumed in proportion to some weight assigned to that device. The weights assigned to the devices can dynamically change, and thus, one aspect of the present invention enables these changes to be tracked and 15 propagated to other file systems needing or desiring this information. Further, recovery of the weights is provided for in the case one or more of the nodes having file systems fail.

20 Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

Brief Description of the Drawings

25 The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the

claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1 depicts one example of a computing environment incorporating and using one or more aspects of the present invention;

FIG. 2 depicts further details of a node of FIG. 1, in accordance with an aspect of the present invention;

FIG. 3a depicts one example of a storage device being partitioned into a plurality of partitions in which each partition is owned by zero or more nodes, in accordance with an aspect of the present invention;

FIG. 3b depicts one example of various statistics associated with each storage device, in accordance with an aspect of the present invention;

FIG. 4 depicts one embodiment of the logic associated with a parallel weighted allocation technique, in accordance with an aspect of the present invention;

FIG. 5 depicts one embodiment of the logic associated with the initialization action of FIG. 4, in accordance with an aspect of the present invention;

5 FIG. 6 depicts one embodiment of the logic associated with the tracking and distribution action of FIG. 4, in accordance with an aspect of the present invention;

10 FIG. 7 depicts one embodiment of the logic associated with the node failure and recovery action of FIG. 4, in accordance with an aspect of the present invention;

15 FIG. 8 depicts one embodiment of the logic associated with the recovery of static weights, in accordance with an aspect of the present invention;

20 FIG. 9 depicts one embodiment of the logic associated with no-state recovery of dynamic weights, in accordance with an aspect of the present invention; and

FIG. 10 depicts one embodiment of the logic associated with full-state recovery of dynamic weights, in accordance with an aspect of the present invention.

Best Mode for Carrying Out the Invention

In accordance with an aspect of the present invention, a plurality of file systems allocate space on one or more storage devices using weights associated with those devices.

5 In particular, the weights associated with the storage devices are used to generate stripe orders, and each stripe order provides to a respective file system the order in which space on individual storage devices is to be allocated and the frequency of allocating space on those devices. The
10 weight associated with each device is distributed to the file systems that are to allocate space on that device, so that the combined allocation remains proportional to the weights. Since the weights can dynamically be adjusted, the various file systems are kept up-to-date of the weight
15 adjustments.

One embodiment of a computing environment incorporating and/or using aspects of the present invention is described with reference to FIG. 1. Computing environment 100 includes one or more nodes 102 (e.g., Node 1,...Node n), which share access to one or more storage devices 104 (e.g., Disk 1...Disk m, or other non-volatile memory). The nodes are coupled to each other and to the storage devices via an interconnect 106. In one example, the interconnect includes a wire connection, a bus, a token ring or a network connection, to name just a few examples. One communications protocol used by one or more of these connections is TCP/IP. It is assumed, in one example, that the nodes do not have shared memory.

00000000000000000000000000000000

As one example, a node 102 includes an operating system 200 (FIG. 2), such as the AIX operating system offered by International Business Machines Corporation. The operating system includes a file system 202 (e.g., a software layer), such as the General Parallel File System (GPFS) offered by International Business Machines Corporation, which is used to manage the allocation of space on various storage devices. In one or more of the embodiments described herein, it is assumed that each node has a single file system, and thus, some of the description references the node. However, in another example, a node may include a plurality of file systems. In that example, each participating file system on the node is kept up-to-date of weight changes and may be involved in recovery.

File system 202 allocates space on various of the storage devices, such that the total allocation on each storage device is proportional to a weight obtained for that device. As used herein, the obtaining of weights can be accomplished in any manner including, but not limited to, receiving the weights, and assigning the weights. The weight obtained for each device is used in determining the allocation policy and allows the file system to balance the allocation across the devices to match individual device capacities and to better utilize the combined throughput of the devices. However, the weights and the allocation policy (i.e., the order and frequency of allocations on each device) are independent of the technique used for the allocation. That is, different allocation techniques can be

used for the allocation. The allocation technique is not tied to the weights. This allows the weights to represent a variety of parameters (e.g., capacity weighting, free space weighting, throughput weighting, round-robin weighting,
5 hybrid weighting, etc., described below), and allows the weights to dynamically change. Thus, the allocation policy can be changed at any time to better suit the current conditions or requirements. Further, any weighting technique used in obtaining the weights need not be known to
10 the allocation technique.

PCT/US2000/0112-US1

Many different allocation techniques can be used to allocate space on the storage devices. Examples of such allocation techniques include a deterministic technique and a randomized technique, each of which is described in detail
15 in co-filed, U.S. Patent Application Serial No. _____, entitled "Determining The Order And Frequency In Which Space Is Allocated On Individual Storage Devices," Sawdon et al., filed _____, which is hereby incorporated herein by reference in its entirety.

20 In a parallel file system, multiple file systems (of one or more nodes) can allocate space on one or more storage devices. As examples, two or more file systems can allocate space on one storage device; and/or two or more file systems can allocate space on two or more storage devices in any
25 combination (e.g., each of a plurality of file systems allocates space on a different device; and/or one or more file systems allocate space on one or more devices.) Any combination of a plurality of file systems allocating space

on one or more devices is possible. Thus, space may be allocated on any one of the storage devices by any one or more of the file systems.

Since a plurality of file systems may allocate space on a particular storage device, in one example, the storage space on a device is partitioned into a plurality of partitions, as depicted in FIG. 3a. As shown in FIG. 3a, a device 300 is partitioned into a plurality of partitions 302a-d, and each partition is owned by zero or more of the nodes. For instance, partition 302a is unowned; partition 302b is owned by Node 1; partition 302c is owned by Node 2; and partition 302d is owned by Nodes 3 and 4. The one or more nodes that own the partition are allowed to allocate space in that partition. (In a further example, ownership could be based on file systems, in which each partition is owned by zero or more file systems, regardless of the nodes in which those file systems reside.)

In one embodiment, ownership information is maintained by a centralized allocation manager. This manager can be a part of one of the nodes participating in the allocation or another node that is used mainly for control and does not actually allocate. Examples of the partitioning of space and of a centralized allocation manager are described in U.S. Patent No. 5,960,446, Schmuck et al., entitled "Parallel File System And Method With Allocation Map," Issued September 28, 1999, which is hereby incorporated herein by reference in its entirety.

The ownership information is maintained as part of various statistics associated with each storage device. In particular, each device has associated therewith the following statistics 310 (FIG. 3b), as one example:

5 (A) Per-Device Total: The device total
represents how much of a particular parameter is
associated with the device. For instance, the total
may indicate the amount of free space on the device.

(B) Per-Partition Information:

10 (1) Owner(s): An indication of the one
or more owners of that particular partition;
and

15 (2) Partition Total: An indication of
how much of the particular parameter is
associated with the partition (e.g., the
amount of free space in the partition).

In accordance with an aspect of the present invention,
each file system that is to allocate space uses a weighted
allocation technique to determine the order in which devices
20 are selected for allocation and the frequency for allocating
space on those devices. The file systems allocating space
on a particular device agree upon the weight for that
device, so that the total allocation of each device remains
proportional to the weight assigned to that device. This

P
O
U
9
-

agreement is provided by contacting the centralized allocation manager, as described below.

The weights used for allocation can dynamically change. Thus, the new values are propagated to the nodes (or file systems) needing or desiring the new information. The tracking and distribution of weights is, therefore, a part of the parallel weighted allocation technique of the present invention. Further, since parallel environments may suffer partial failures, with one or more nodes failing and restarting independently, the allocation technique of the present invention also includes recovery.

One example of a weighted allocation technique of a parallel file system is described with reference to FIG. 4. As shown in FIG. 4, the technique includes three main actions, including initialization, STEP 400, tracking and distribution of weights, STEP 402, and node failure and recovery, STEP 404. Each of these actions can be implemented in various ways. Two possible embodiments for each action are described herein. The first embodiment is referred to as a no-state embodiment, which uses a minimal amount of state, but has a higher time for recovery from a node failure. The second embodiment is referred to as a full-state embodiment, in which the allocation manager is used to maintain the partition ownership information, as well as complete per-device-per-partition counters. This extra state serves to reduce the time for recovery. Each of these embodiments for each of the actions is described in further detail below.

One example of the initialization action is described in further detail with reference to FIG. 5. Both embodiments of this action (i.e., the no-state embodiment and the full-state embodiment) perform the actions depicted
5 in FIG. 5.

Initially, the file system selects an allocation manager, STEP 500. In one example, the first node that attempts to run the initialization logic is designated as the allocation manager. The other nodes are referred to as
10 client nodes. The client nodes locate the allocation manager using, for instance, a global naming service, and wait for the allocation manager's initialization to complete.

Subsequent to appointing the allocation manager, the
15 allocation manager determines the initial weights to be used for allocation, STEP 502. The allocation manager may determine the weights serially working alone or in parallel by enlisting the assistance of one or more of the client nodes.

20 The initial weights depend on the weighting technique used. A variety of weighting techniques are available including techniques based on static parameters, as well as techniques based on dynamic parameters. Examples of various techniques include, for instance, the following:

(1) Round-Robin Weighting - To implement a simple round-robin allocation, the weight of each device is set to 1. Using an equal weight for each device, the technique will allocate space on each device an equal number of times.

(2) Capacity Weighting - To better distribute the allocations across uneven sized devices, the weights can be assigned using the relative capacity of each device. This weighting technique causes the devices to fill in the same proportion (i.e., the percentage utilized on each device is the same, regardless of the capacity of the device). Consequently, the expected I/O load on each device is also in proportion to the device's capacity.

For capacity weighting, the allocation manager determines the maximum storage capacity of each device. This can be done in a number of ways, such as examining a descriptor for each device.

(3) Free Space Weighting - In this dynamic weighting technique, the weights may be based upon the relative amount of free space on each device. Under this technique, devices with a higher percentage of free space receive proportionately more allocations. This serves to rebalance unevenly filled devices, which may have resulted from adding new devices to an existing system or previously using round-robin allocation on uneven sized

devices. The weights can be adjusted dynamically to account for changes in the relative amount of free space. For devices that are evenly filled, the free space weighting technique is equivalent to capacity weighting.

5

For free space weighting, each device is examined to determine the number of free blocks on each device. In one example, the allocation map of the device can provide this information.

10

(Allocation maps are described in U.S. Patent No. 5,960,446, Schmuck et al., entitled "Parallel File System And Method With Allocation Map," Issued September 28, 1999, which is hereby incorporated herein by reference in its entirety.)

15

(4) Throughput Weighting - The weights can also be assigned based on the relative performance of each device. Devices with higher throughput receive proportionately more allocations and consequently, more I/O requests on the average. This weighting attempts to maximize the total throughput of the combined devices.

20

There are a number of ways to determine the throughput weights during initialization, including, for instance, reading the device throughput from a table or by measuring the device throughput by measuring the actual throughput to each device, while the system is under a maximal I/O load.

25

(5) Hybrid Weighting - Not only can the weights be changed dynamically, the technique for assigning the weights can also be changed. Furthermore, a combination of two or more weighting techniques to produce a hybrid weighting can be used. This may be accomplished by computing the normalized weight for each device under more than one technique, then adding the normalized weights for a device from each desired technique. This allows the system to tailor the allocation to the current requirements and to change the allocation as the system changes.

Continuing with reference to FIG. 5, after the initial weights have been collected, the allocation manager propagates the weights to the other nodes (or other file systems), STEP 504. In one example, the weights are propagated to all of the nodes participating in allocation. In another example, the weight of a particular device is only propagated to the nodes that are to use that weight. The propagation can be performed using a number of techniques, including passing messages or broadcasting.

Thereafter, each file system that is to allocate uses the weights to initialize a local weighted allocation technique, STEP 506. The local weighted allocation technique is a technique executed by the file system to generate the stripe order used to define the order and frequency of allocation on the storage devices. This technique includes, for instance, an initialization step

that normalizes the weights and sets some variables; and a stripe order generation step that uses the normalized weights to determine the stripe order. Other steps may also be included, depending on the type of allocation technique.

5 Various weighted allocation techniques are described in detail in co-filed, U.S. Patent Application Serial No. _____, entitled "Determining The Order And Frequency In Which Space Is Allocated On Individual Storage Devices," Sawdon et al., filed _____, which is hereby
10 incorporated herein by reference in its entirety. Examples of these techniques include a deterministic technique and a randomized technique. In one example, if the deterministic technique is used, the starting position within the stripe order is random. Thus, different file systems may begin at
15 different positions within the stripe orders.

Initializing the local weighted allocation technique completes the initialization action. As mentioned above, both the no-state and full-state embodiments perform the above-described actions. However, in addition to the above,
20 the full-state embodiment, when using dynamic weighting, like free space weighting, also saves the free space for each device in each partition in a per-device-per-partition table (See FIG. 3b).

Referring back to FIG. 4, subsequent to performing the
25 initialization, the action of tracking and distribution of weights is performed, STEP 402. Weights based on dynamic information, such as free space per device, is periodically

updated in order to maintain their accuracy. Weights based on static information, such as capacity, is updated when, for instance, the configuration changes or when there is a change in the allocation policy. One embodiment of the logic employed in tracking and distributing weights is described with reference to FIG. 6. This particular example is described with reference to the tracking and distribution of free space (a dynamic weight). However, the logic is similarly applicable to other dynamic weights or to static weights.

Referring to FIG. 6, each of various nodes tracks the changes in information (i.e., dynamic information and/or static information), STEP 600. As one example, for free space weighting, each appropriate node tracks the number of allocations and deallocations that it performs on each device. The net allocations per device, called the delta, is the difference in free space on each device caused by operations at that node. The client node accumulates the deltas until some threshold (e.g., 100 operations) is met. When the threshold is met or at another predefined event (e.g., every 30 seconds), the node informs the allocation manager of the changes, STEP 602. In particular, a communications mechanism is used by the client node to send the deltas to the allocation manager. After successfully sending the deltas, the client node then resets its delta counters to zero.

Upon receiving the deltas from a client, the allocation manager adds them to the total free space counters, which

are maintained by the allocation manager (see per-device total 312 of FIG. 3b), STEP 604. That is, the allocation manager adjusts the weights for particular devices based on the deltas that it receives.

5 Subsequent to adjusting the weights, the adjusted weights (e.g., the new total free space counters for the devices) are returned to the client in reply to the message sending the deltas, STEP 606. As the client receives the adjusted weights, the client reinitializes its local
10 allocation technique using the adjusted weights, STEP 608. Thus, a new stripe order is generated.

15 The above-described communication between the clients and allocation manager enable the clients, once they have communicated with the allocation manager, to agree on the weights to be used. This agreement can come quicker, if the nodes are informed quicker about the new weights. The decision of when and how to tell the clients is embodiment dependent.

20 For example, the threshold used by the clients for sending the deltas also serves to bound the difference between the total free space counters maintained by the allocation manager and the actual amount of free space on each device. To maintain the same degree of accuracy of the weights used by the client nodes, any large change in the
25 weights caused by, for instance, a change in the allocation policy, a change in the hardware configuration, or a large accumulative change due to deltas received from a set of

very active client nodes, causes the allocation manager, in one example, to immediately send the new weights to the client nodes. This change in weights does not effect the current deltas stored at the nodes.

5 The above actions are performed for each of the two embodiments described herein (i.e., the no-state embodiment and the full-state embodiment). However, for the full-state embodiment, each client node maintains separate delta counters for each partition that it modifies. Upon
10 receiving the per-partition deltas, the allocation manager updates the per-device-per-partition counters 318 (FIG. 3b), as well as the device totals 312.

15 Returning to FIG. 4, in addition to the tracking and distribution of weights, which enables the rebalancing of a system based on weighted allocation, the parallel weighted allocation technique of the present invention also provides for recovery from a node failure, STEP 404. Nodes in a parallel file environment may fail or be restarted independently of each other. To handle node failures, the
20 volatile state lost by the failed node is to be reconstructed by another node. This recovery depends on a number of factors, including, for instance: whether the failed node is a client or acting as the allocation manager; on whether the weights are static or dynamic; and for
25 dynamic weights, it also depends on the amount of state maintained by the allocation manager.

00000000000000000000000000000000

There are three main cases to be considered herein:
recovery using static weights, recovery using dynamic
weights with a no-state embodiment, and recovery using
dynamic weights with a full-state embodiment. For each of
5 these cases, an example technique for recovery is described
for failure from a single client node or from the allocation
manager's node. The techniques are easily extended to handle
multi-node failures, as long as a quorum of the nodes remain
available, in one example. The recovery from node failure,
10 which is managed by the file system, is further described
with reference to FIGs. 7-10.

Referring to FIG. 7, initially a determination is made
as to the type of recovery that is needed, STEP 700. For
instance, a determination is made as to whether recovery of
15 static weights is needed, INQUIRY 702. If it is a recovery
of static weights, then processing continues with the logic
of FIG. 8, STEP 704.

Referring to FIG. 8, initially a determination is made
as to whether a client node failed, INQUIRY 800. If a
20 client node failed, then no additional recovery is needed,
STEP 802. However, if it was not a client node, and
therefore, an allocation manager that failed, then the
static weights are recovered, STEP 804. In one example, the
static weights are recovered by obtaining them from a client
25 node or reconstructing them from other available
information. The nodes that did not fail can continue
allocating as usual, even throughout the recovery of the
failed node.

P009-2000-0112-US1

Referring back to FIG. 7, if the recovery is not of static weights, then it is assumed to be recovery of dynamic weights. Therefore, a determination is made as to whether it is recovery of dynamic weights with no-state, INQUIRY

5 706. If it is a no-state recovery of dynamic weights, then recovery proceeds as described with reference to FIG. 9, STEP 708. Again, the examples are described with reference to free space, but can be extended to other dynamic weights.

10 Referring to FIG. 9, initially a determination is made as to whether it was a client node that failed, INQUIRY 900. If the client node failed, then the allocation manager checks the partition ownership information for partitions that are not owned and marks these partitions as unavailable to prevent them from being assigned to a client node until the recovery associated with the partition is complete, STEP 902. (When a node fails, partitions owned by that node become unowned.)

15 Additionally, the allocation manager checks the partition ownership information for partitions owned by more than one node. For each shared partition, it sends a revoke-ownership message to all the owners except one, STEP 904. This minimizes the number of nodes to be involved in the recovery.

20 The allocation manager then sets the per-device free space totals to zero, STEP 906, and sends a broadcast message to the non-failed nodes asking them for the

per-device free space counts for the partitions that are owned by that node, STEP 908.

Upon receiving this message, each appropriate client node stops allocating and resets its delta counters to zero.

5 Further, it returns the per-device free space count for each owned partition to the allocation manager. The node may then resume allocating space in the partitions that it currently owns.

As the allocation manager receives the replies, STEP 910, the per-device free space counts are added to the totals, STEP 912. Further, the free space in all unowned partitions is also recovered, STEP 914. This may be done serially by the allocation manager or in parallel by enlisting the aid of the client nodes. In one example, 15 since the no-state embodiment lacks the state information to delimit the recovery to only the partitions modified by the failed node, recovery of a failed node includes the reading of the non-volatile allocation maps in order to reconstruct the per-device free space totals. As each unknown partition 20 is recovered, it becomes eligible for assignment and is marked as available. This completes the no-state recovery of dynamic weights for a failed client. Upon completion of the recovery, a value for the dynamic weight (e.g., total free space) has been recomputed, and this adjusted weight 25 can be forwarded to one or more file systems, as described above.

Returning to INQUIRY 900, if the failed node is the allocation manager, then a new allocation manager is selected, STEP 918. In one example, the new allocation manager is selected by assigning the function to the non-failed node with the lowest id/address.

The newly assigned allocation manager rebuilds the partition ownership information, STEP 920. In one example, this is accomplished by requesting information from the other nodes regarding the partitions that they own. For example, the allocation manager sends a broadcast message to the surviving nodes asking them to identify the partitions that they own. Thereafter, recovery proceeds with STEP 902, as described above.

Returning to FIG. 7, if recovery is of dynamic weights with full-state capabilities, INQUIRY 706, then recovery proceeds as described with reference to FIG. 10, STEP 710. In this example, the state maintained by the full-state embodiment enables the recovery from a failed node to be of the partitions owned by the node that failed. It also allows the non-failed client nodes to continue allocating through the recovery.

Referring to FIG. 10, initially, a determination is made as to whether it was a client node that failed, INQUIRY 1000. If it is a client node that failed, the allocation manager checks the partition ownership information for partitions that were owned by the failed node. These partitions are marked as unavailable to prevent them from

DISCRETE STATEMENT

being reassigned until after their recovery is complete,
STEP 1002.

The allocation manager then checks the partition
ownership information for partitions owned by the failed
node and shared with one or more nodes. For each such
shared partition, the allocation manager sends a
revoke-ownership message to all non-failed owners, STEP
1004. Upon receiving this message, a client releases
ownership on the partition and sets the partition's delta
counters to zero.

Thereafter, the free space in the unavailable
partitions is recovered either serially by the allocation
manager or in parallel by enlisting the aid of one or more
of the client nodes, STEP 1006. As each partition is
recovered, the per-device totals and
per-device-per-partition information is updated and the
partition is marked as available for assignment. This
completes the recovery from a failed client node.

Returning to INQUIRY 1000, if it was the allocation
manager that failed, then recovery proceeds as follows.
Initially, a new allocation manager is selected, STEP 1008.
In one example, this is accomplished by assigning the
function to the non-failed node with the lowest id/address.

The newly assigned allocation manager rebuilds the
partition ownership information, STEP 1010. In one example,
the information is built by sending a broadcast message to

the surviving nodes asking them for the partitions that they own. Partitions that are unowned are marked as unavailable by the allocation manager to prevent them from being allocated until recovery is complete, STEP 1012.

5 The allocation manager then checks the partition ownership information for partitions owned by one or more nodes. For each shared partition, the allocation manager sends a revoke-ownership message to all the owners except one, STEP 1014. Upon receiving this message, a client
10 releases ownership on the partition and sets the partition's delta counters to zero.

The allocation manager then sends a broadcast message to the nodes asking them to send the per-device free space information for each partition that they own, STEP 1016.

15 Upon receiving this message, a client resets the partition's delta counters to zero and returns the per-device free space information to the allocation manager.

20 As the allocation manager receives the replies, it updates the per-device-per-partition information, as well as the per-device totals, STEP 1018.

25 Subsequently, the free space in the unavailable partitions is recovered either serially by the allocation manager or in parallel by enlisting the aid of one or more of the client nodes, STEP 1020. As each partition is recovered, the per-device totals and per-device-per-partition information is updated and the partition is marked

as available for assignment. This completes the full-state recovery of dynamic weights.

In accordance with an aspect of the present invention, the recovery techniques (both for static and dynamic weights) maintain goal weight values associated with each device. In particular, in one example, each device has a goal weight associated therewith. In the case of static weighting, the goal weights are equivalent to the static weights, and thus, no distinction need be made. However, in dynamic weighting, the goal weights (which are static, in one example) may be different than the weights being used, at any particular time. That is, the weights being used may have been adjusted such that the goal weights are satisfied. This maintaining of the goal weights of the devices is accomplished even if one or more nodes (or file systems) fail. Also, it is maintained even if one or more storage devices fail and are restarted or replaced.

Described in detail above is a parallel weighted allocation capability that enables a plurality of file systems to use weighted allocation to allocate space on one or more storage devices. The space is allocated on the shared storage devices, such that the space on each device is consumed in proportion to some weight assigned to that device. This allows the allocation to be balanced across the devices, such that the load on each device is proportional to the weight assigned to that device. For a parallel environment, the weight assigned to each device is distributed to the various file systems using that weight,

so that the combined allocation remains proportional to the weights. Furthermore, the file systems are kept up-to-date as the weights change or the technique of assigning the weights change.

5 In one example, different file systems can use
different allocation techniques in order to allocate space
on various storage devices. In one example, this is
accomplished by grouping the storage devices into groups,
wherein a file system using one allocation technique
10 allocates space on one group of devices and a file system
using a different technique allocates space on a different
group of devices. In another example however, the grouping
is not used.

15 Further, the weights obtained by various file systems
can represent different parameters. For instance, one file
system can obtain weights based on free space and another
file system can obtain weights based on capacity. Again, in
one example, the storage devices may be grouped into
differing groups in order to accommodate the usage of
20 different weighting techniques. In another example however,
the grouping is not used.

25 In yet a further example, the allocation policy can be
set at various levels. In particular, one or more stripe
orders can be generated and used to allocate space across
the storage devices. As examples, one stripe order can be
used for all allocations of a file system; in another
example, the storage devices are partitioned into groups,

and a stripe order is generated and used for each group; and in yet another embodiment, a stripe order is generated for each file that is going to have space allocated therefor. Thus, in the last example, one stripe order can be used to 5 allocate space for one file, and another stripe order (the same or different) can be used to allocate space for another file. In any of the above scenarios, the stripe orders are generated, as described above.

Although weighted allocation depends on the weights to 10 determine the allocation policy, the allocation techniques themselves are independent of the actual weights assigned to each device. The weights can be changed at any time, to adjust the load on each device, as needed or desired. Furthermore, the technique of assigning the weights can be 15 changed at any time. This allows the allocation policy to be set dynamically and adjusted to meet the current requirements of the system. Further, the changing of the allocation policy can occur without restarting the file system.

The weights assigned to the devices can be dynamically 20 changed to represent different values and/or to represent a different operating parameter (e.g., capacity, free space, I/O throughput, round-robin, hybrid). Further, the weighting assignment technique need not be known to the 25 allocation technique. Further, the allocation technique can accommodate various data streams, including video streams and general data streams. This is because the allocation technique does not know and need not know apriori the length

of the data streams and/or the access patterns of those data streams.

The allocation capability of the present invention is also able to stripe according to weight across a plurality of heterogeneous storage devices. That is, the storage devices may be of different sizes, different capacities and/or of different speeds. These heterogeneous devices can be utilized and that utilization can be maximized. For instance, storage usage can be maximized and/or throughput can be maximized.

Additionally, the allocation capability of the present invention can automatically compensate for an imbalance in the parallel file environment. Such an imbalance can be caused by adding devices to the system, removing devices from the system, or for any other reasons. The rebalancing of the environment is performed without necessarily restriping space already striped. In one example, the rebalancing is accomplished by obtaining new, different and/or additional weights and using an allocation technique to allocate space based on those weights.

The above-described computing environment is offered as only one example. One or more aspects of the present invention can be incorporated and used with many types of computing units, computers, processors, nodes, systems, work stations and/or environments without departing from the spirit of the present invention.

Various of the embodiments described above reference a node receiving information, providing information or performing some task. If, however, the node includes a plurality of file systems, then one or more of those file systems on the node may perform those actions.

The present invention can be included in an article of manufacture (e.g., one or more computer program products) having, for instance, computer usable media. The media has embodied therein, for instance, computer readable program code means for providing and facilitating the capabilities of the present invention. The article of manufacture can be included as a part of a computer system or sold separately.

Additionally, at least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform the capabilities of the present invention can be provided.

The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications,

additions, substitutions and the like can be made without departing from the spirit of the invention and these are therefore considered to be within the scope of the invention as defined in the following claims.

5

00000000000000000000000000000000

Claims

What is claimed is:

1. A method of managing the allocation of space on storage devices of a computing environment, said method comprising:

obtaining one or more weights for one or more storage devices of said computing environment; and

allocating space on at least one storage device of said one or more storage devices in proportion to at least one weight obtained for the at least one storage device, wherein said allocating is performed by a plurality of file systems of said computing environment.

15 2. The method of claim 1, wherein each of said plurality of file systems is located on a separate node of said computing environment.

20 3. The method of claim 1, wherein said plurality of file systems are located on one or more nodes of said computing environment.

100-1234567890

4. The method of claim 1, wherein said allocating comprises executing an allocation technique by each file system of said plurality of file systems, wherein at least one file system of said plurality of file systems is running
5 a different allocation technique than at least one other file system of said plurality of file systems.

10 5. The method of claim 1, wherein each storage device of said at least one storage device is partitioned into a plurality of partitions, and wherein one or more partitions of each storage device are owned by one or more file systems of said plurality of file systems.

15 6. The method of claim 1, wherein said allocating comprises allocating space on a plurality of storage devices by a plurality of file systems, wherein each file system of said plurality of file systems allocates space on one or more storage devices of said plurality of storage devices.

7. The method of claim 1, wherein said obtaining comprises using at least an allocation manager to obtain said one or more weights.

20 8. The method of claim 7, wherein said using comprises using said allocation manager and at least one node of said computing environment to obtain said one or more weights.

25 9. The method of claim 1, wherein said one or more weights represent at least one parameter of said computing environment.

100-2000-0112-US1

10. The method of claim 1, wherein said allocating is independent of the obtaining of said one or more weights, wherein the allocating need not have knowledge of at least one of what the weights represent and how the weights were obtained.

5

11. The method of claim 1, wherein at least one storage device of said one or more storage devices has one or more different characteristics than at least one other storage device of said one or more storage devices.

10

12. The method of claim 1, further comprising propagating the at least one weight to at least one file system of said plurality of file systems.

15

13. The method of claim 1, further comprising:

tracking changes associated with at least one weight of said one or more weights;

adjusting said at least one weight based on the tracked changes; and

20

propagating the at least one adjusted weight to a file system of said computing environment, wherein said at least one adjusted weight is usable in allocating space on at least one storage device.

14. The method of claim 13, wherein said tracking is performed by the file system.

15. The method of claim 13, wherein said tracking is performed by a plurality of file systems, and wherein said propagating comprises propagating the at least one adjusted weight to the plurality of file systems that performed the tracking.

16. The method of claim 13, further comprising informing an allocation manager, at a predefined event, of the tracked changes, and wherein said allocation manager performs the adjusting and the propagating.

17. The method of claim 1, further comprising informing said plurality of file systems of changes in said at least one weight, wherein said changes are usable in further allocating space.

18. The method of claim 1, further comprising adjusting at least one weight of said one or more weights, in response to a failure of a file system of said computing environment.

19. The method of claim 18, wherein said adjusting comprises at least one of:

using information provided by at least one other file system of said computing environment to adjust said at least one weight; and

5

using information obtained from reading at least one storage device associated with said at least one weight to adjust said at least one weight.

10 20. The method of claim 1, further comprising maintaining at least one weight of said one or more weights, in response to a failure of a file system of said computing environment.

15 21. The method of claim 1, wherein one file system of said plurality of file systems allocates space on said at least one storage device for a given file, and wherein said allocating for that given file is based on an allocation policy that uses said at least one weight.

20 22. The method of claim 21, wherein said one file system allocates space on one or more storage devices for another file, and wherein the allocating for that another file is based on another allocation policy that uses one or more weights associated with the one or more storage devices.

25

23. A method of managing the allocation of space on storage devices of a computing environment, said method comprising:

5 obtaining a weight for each storage device of at least a subset of storage devices of a plurality of storage devices of said computing environment; and

10 allocating space on each storage device of said at least a subset of storage devices in proportion to the weight assigned to the storage device, wherein said allocating is performed by a plurality of file systems, such that each file system of said plurality of file systems allocates space on one or more storage devices of said at least said subset of storage devices.

15

24. A system of managing the allocation of space on storage devices of a computing environment, said system comprising:

5 means for obtaining one or more weights for one or more storage devices of said computing environment; and

10 means for allocating space, by a plurality of file systems of said computing environment, on at least one storage device of said one or more storage devices in proportion to at least one weight obtained for the at least one storage device.

15 25. The system of claim 24, wherein each of said plurality of file systems is located on a separate node of said computing environment.

26. The system of claim 24, wherein said plurality of file systems are located on one or more nodes of said computing environment.

20 27. The system of claim 24, wherein said means for allocating comprises means for executing an allocation technique by each file system of said plurality of file systems, wherein at least one file system of said plurality of file systems is running a different allocation technique than at least one other file system of said plurality of file systems.

100-1000

28. The system of claim 24, wherein each storage device of said at least one storage device is partitioned into a plurality of partitions, and wherein one or more partitions of each storage device are owned by one or more file systems of said plurality of file systems.

29. The system of claim 24, wherein said means for allocating comprises means for allocating space on a plurality of storage devices by a plurality of file systems, wherein each file system of said plurality of file systems allocates space on one or more storage devices of said plurality of storage devices.

30. The system of claim 24, wherein said means for obtaining comprises means for using at least an allocation manager to obtain said one or more weights.

15 31. The system of claim 30, wherein said means for using comprises means for using said allocation manager and at least one node of said computing environment to obtain said one or more weights.

20 32. The system of claim 24, wherein said one or more weights represent at least one parameter of said computing environment.

00000000000000000000000000000000

33. The system of claim 24, wherein said means for allocating is independent of the means of obtaining of said one or more weights, wherein the means for allocating need not have knowledge of at least one of what the weights represent and how the weights were obtained.

5

34. The system of claim 24, wherein at least one storage device of said one or more storage devices has one or more different characteristics than at least one other storage device of said one or more storage devices.

10

35. The system of claim 24, further comprising means for propagating the at least one weight to at least one file system of said plurality of file systems.

36. The system of claim 24, further comprising:

15

means for tracking changes associated with at least one weight of said one or more weights;

means for adjusting said at least one weight based on the tracked changes; and

20

means for propagating the at least one adjusted weight to a file system of said computing environment, wherein said at least one adjusted weight is usable in allocating space on at least one storage device.

37. The system of claim 36, wherein said means for tracking comprises means for tracking by the file system.

38. The system of claim 36, wherein said means for tracking comprises means for tracking by a plurality of file systems, and wherein said means for propagating comprises means for propagating the at least one adjusted weight to the plurality of file systems used in the tracking.

39. The system of claim 36, further comprising means for informing an allocation manager, at a predefined event, of the tracked changes, and wherein said allocation manager performs the adjusting and the propagating.

40. The system of claim 24, further comprising means for informing said plurality of file systems of changes in said at least one weight, wherein said changes are usable in further allocating space.

41. The system of claim 24, further comprising means for adjusting at least one weight of said one or more weights, in response to a failure of a file system of said computing environment.

42. The system of claim 41, wherein said means for adjusting comprises at least one of:

means for using information provided by at least one other file system of said computing environment to adjust said at least one weight;
5 and

means for using information obtained from reading at least one storage device associated with said at least one weight to adjust said at least one weight.
10

43. The system of claim 24, further comprising means for maintaining at least one weight of said one or more weights, in response to a failure of a file system of said computing environment.

15 44. The system of claim 24, wherein one file system of said plurality of file systems allocates space on said at least one storage device for a given file, and wherein the allocating for that given file is based on an allocation policy that uses said at least one weight.

20 45. The system of claim 44, wherein said one file system allocates space on one or more storage devices for another file, and wherein the allocating for that another file is based on another allocation policy that uses one or more weights associated with the one or more storage
25 devices.

46. A system of managing the allocation of space on storage devices of a computing environment, said system comprising:

means for obtaining a weight for each storage device of at least a subset of storage devices of a plurality of storage devices of said computing environment; and

a plurality of file systems adapted to allocate space on each storage device of said at least a subset of storage devices in proportion to the weight assigned to the storage device, wherein each file system of said plurality of file systems allocates space on one or more storage devices of said at least said subset of storage devices.

DELETION REQUESTED

47. A system of managing the allocation of space on storage devices of a computing environment, said system comprising:

5 at least one node adapted to obtain one or more weights for one or more storage devices of said computing environment; and

10 a plurality of nodes adapted to allocate space on at least one storage device of said one or more storage devices in proportion to at least one weight obtained for the at least one storage device.

48. The system of claim 47, wherein said plurality of nodes comprise said at least one node.

PCT/US2009/046050

49. At least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform a method of managing the allocation of space on storage devices of a computing environment, said method comprising:

5 obtaining one or more weights for one or more storage devices of said computing environment; and

10 allocating space on at least one storage device of said one or more storage devices in proportion to at least one weight obtained for the at least one storage device, wherein said allocating is performed by a plurality of file systems of said computing environment.

15 50. The at least one program storage device of claim 49, wherein each of said plurality of file systems is located on a separate node of said computing environment.

20 51. The at least one program storage device of claim 49, wherein said plurality of file systems are located on one or more nodes of said computing environment.

52. The at least one program storage device of claim
49, wherein said allocating comprises executing an
allocation technique by each file system of said plurality
of file systems, wherein at least one file system of said
plurality of file systems is running a different allocation
technique than at least one other file system of said
plurality of file systems.

53. The at least one program storage device of claim
49, wherein each storage device of said at least one storage
device is partitioned into a plurality of partitions, and
wherein one or more partitions of each storage device are
owned by one or more file systems of said plurality of file
systems.

54. The at least one program storage device of claim
49, wherein said allocating comprises allocating space on a
plurality of storage devices by a plurality of file systems,
wherein each file system of said plurality of file systems
allocates space on one or more storage devices of said
plurality of storage devices.

20 55. The at least one program storage device of claim
49, wherein said obtaining comprises using at least an
allocation manager to obtain said one or more weights.

25 56. The at least one program storage device of claim
55, wherein said using comprises using said allocation
manager and at least one node of said computing environment
to obtain said one or more weights.

57. The at least one program storage device of claim
49, wherein said one or more weights represent at least one
parameter of said computing environment.

58. The at least one program storage device of claim
5 49, wherein said allocating is independent of the obtaining
of said one or more weights, wherein the allocating need not
have knowledge of at least one of what the weights represent
and how the weights were obtained.

59. The at least one program storage device of claim
10 49, wherein at least one storage device of said one or more
storage devices has one or more different characteristics
than at least one other storage device of said one or more
storage devices.

60. The at least one program storage device of claim
15 49, wherein said method further comprises propagating the at
least one weight to at least one file system of said
plurality of file systems.

61. The at least one program storage device of claim
49, wherein said method further comprises:

tracking changes associated with at least one
weight of said one or more weights;

5 adjusting said at least one weight based on
the tracked changes; and

10 propagating the at least one adjusted weight
to a file system of said computing environment,
wherein said at least one adjusted weight is
usable in allocating space on at least one
storage device.

62. The at least one program storage device of claim
61, wherein said tracking is performed by the file system.

15 63. The at least one program storage device of claim
61, wherein said tracking is performed by a plurality of
file systems, and wherein said propagating comprises
propagating the at least one adjusted weight to the
plurality of file systems that performed the tracking.

20 64. The at least one program storage device of claim
61, wherein said method further comprises informing an
allocation manager, at a predefined event, of the tracked
changes, and wherein said allocation manager performs the
adjusting and the propagating.

65. The at least one program storage device of claim
49, wherein said method further comprises informing said
plurality of file systems of changes in said at least one
weight, wherein said changes are usable in further
5 allocating space.

66. The at least one program storage device of claim
49, wherein said method further comprises adjusting at least
one weight of said one or more weights, in response to a
failure of a file system of said computing environment.

10 67. The at least one program storage device of claim
66, wherein said adjusting comprises at least one of:

using information provided by at least one
other file system of said computing environment
to adjust said at least one weight; and

15 using information obtained from reading at
least one storage device associated with said
at least one weight to adjust said at least one
weight.

68. The at least one program storage device of claim
20 49, wherein said method further comprises maintaining at
least one weight of said one or more weights, in response to
a failure of a file system of said computing environment.

69. The at least one program storage device of claim
49, wherein one file system of said plurality of file
systems allocates space on said at least one storage device
for a given file, and wherein the allocating for that given
5 file is based on an allocation policy that uses said at
least one weight.

70. The at least one program storage device of claim
69, wherein said one file system allocates space on one or
more storage devices for another file, and wherein the
10 allocating for that another file is based on another
allocation policy that uses one or more weights associated
with the one or more storage devices.

5

71. At least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform a method of managing the allocation of space on storage devices of a computing environment, said method comprising:

obtaining a weight for each storage device of at least a subset of storage devices of a plurality of storage devices of said computing environment; and

10 allocating space on each storage device of said at least a subset of storage devices in proportion to the weight assigned to the storage device, wherein said allocating is performed by a plurality of file systems, such that each file system of said plurality of file systems allocates space on one or more storage devices of said at least said subset of storage devices.

15

* * * * *

**A PLURALITY OF FILE SYSTEMS USING
WEIGHTED ALLOCATION TO ALLOCATE SPACE
ON ONE OR MORE STORAGE DEVICES**

Abstract of the Disclosure

5 Space is allocated on storage devices in proportion to
weights associated with the storage devices. The space is
allocated by a plurality of file systems. In particular,
space may be allocated on any one of the devices by one or
more of the file systems. The weights can be dynamically
10 adjusted at any time in order to accommodate changes in the
system and to better utilize the storage devices. However,
since more than one file system may be allocating space on
one or more of the storage devices, changes in the weights
are propagated to the various file systems that may utilize
15 the information.

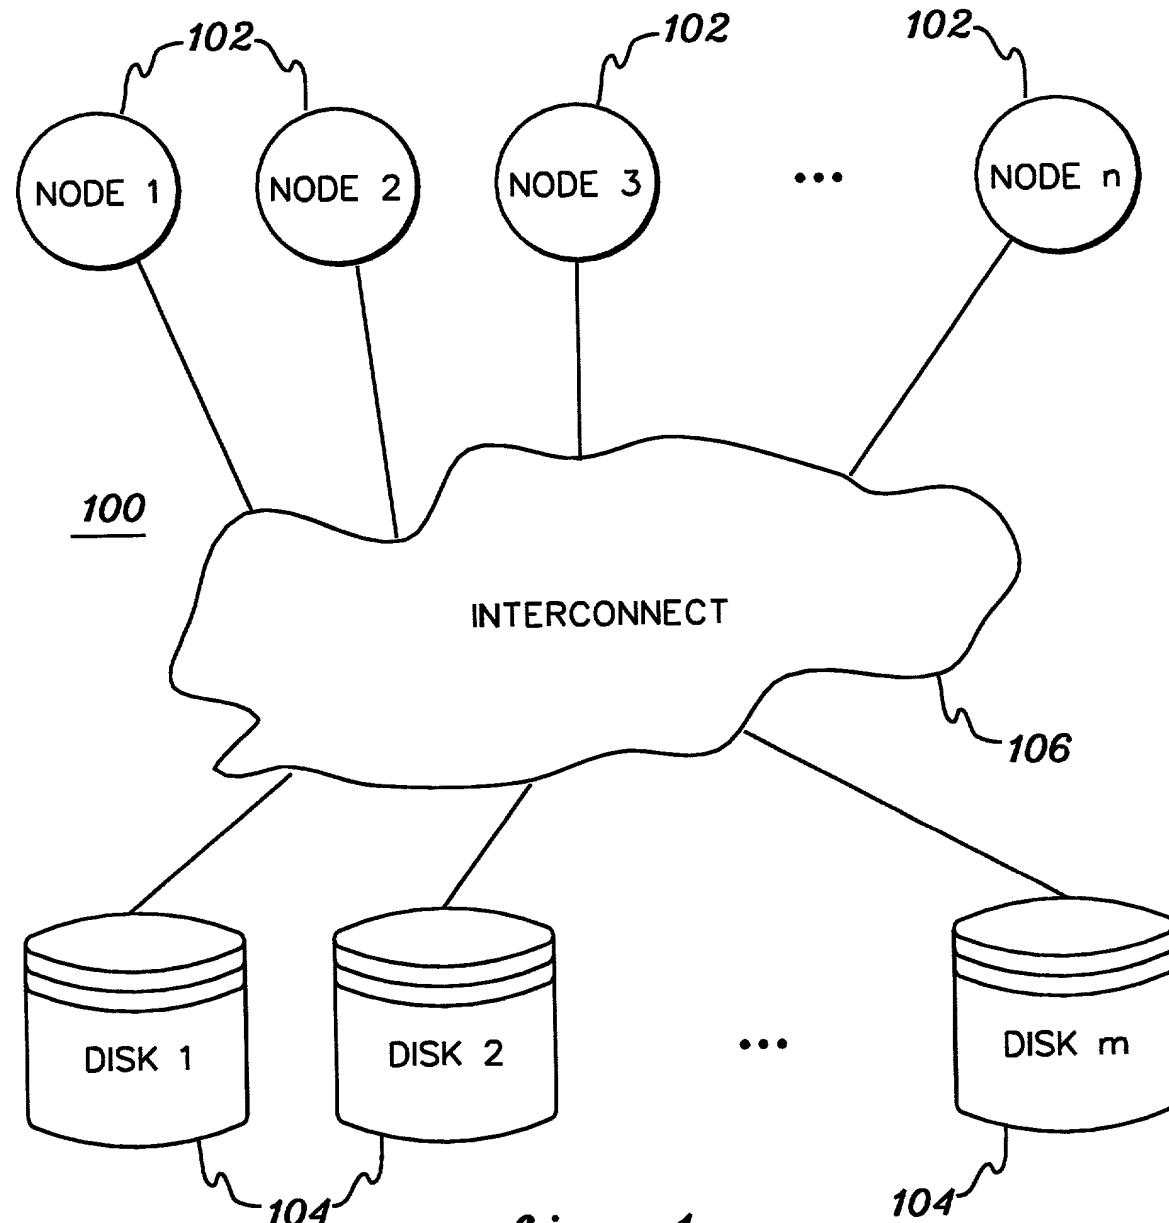


fig. 1

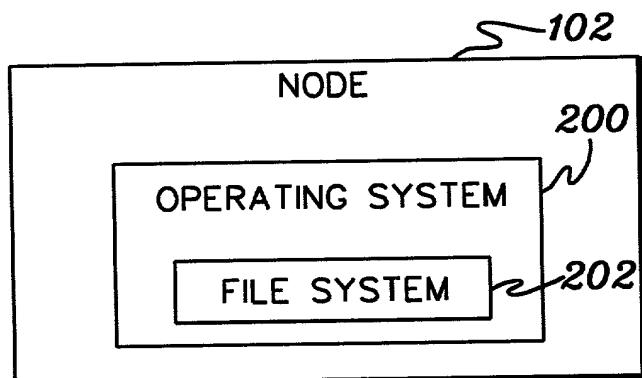


fig. 2

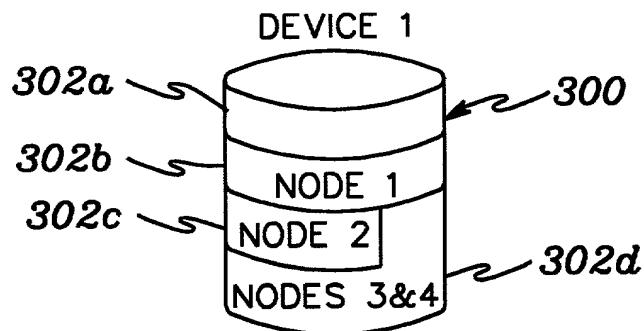


fig. 3A

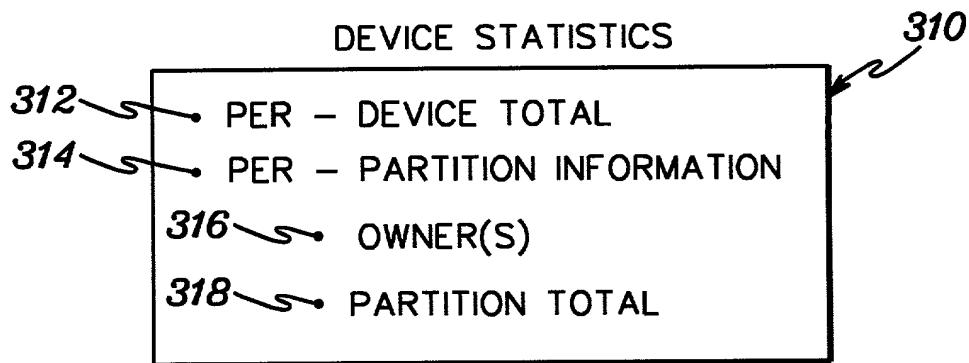


fig. 3B

PARALLEL WEIGHTED ALLOCATION

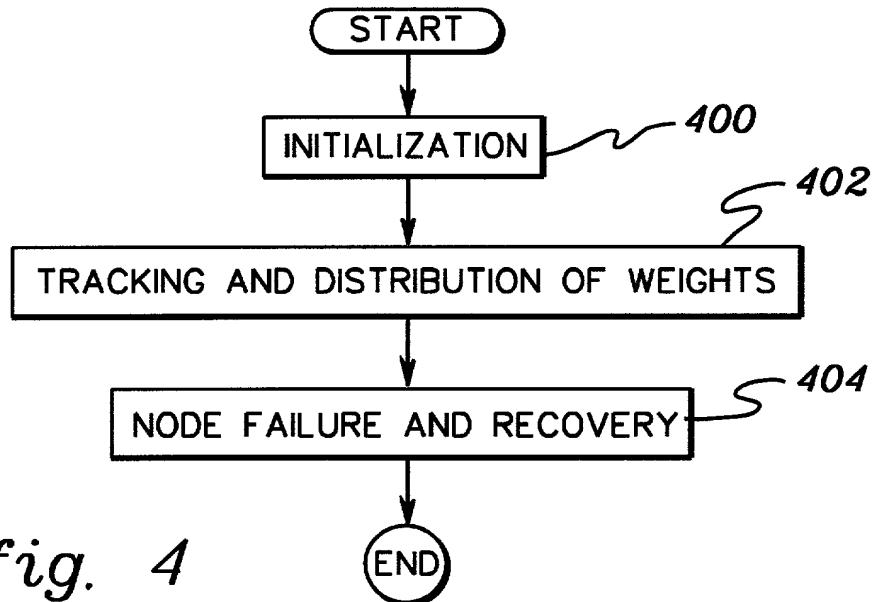


fig. 4

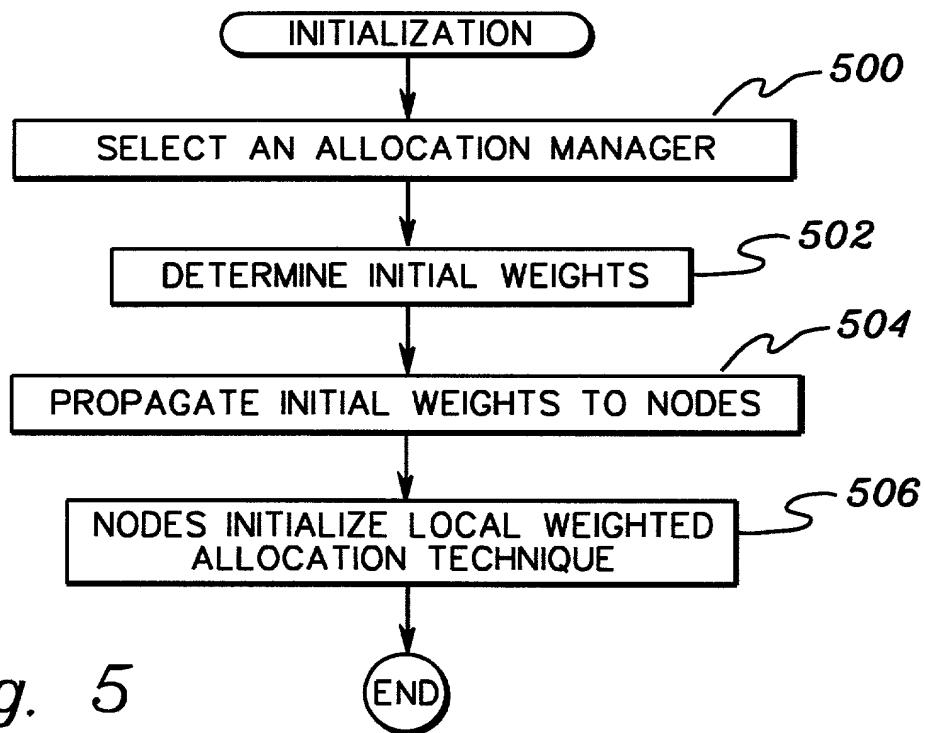


fig. 5

00000000000000000000000000000000

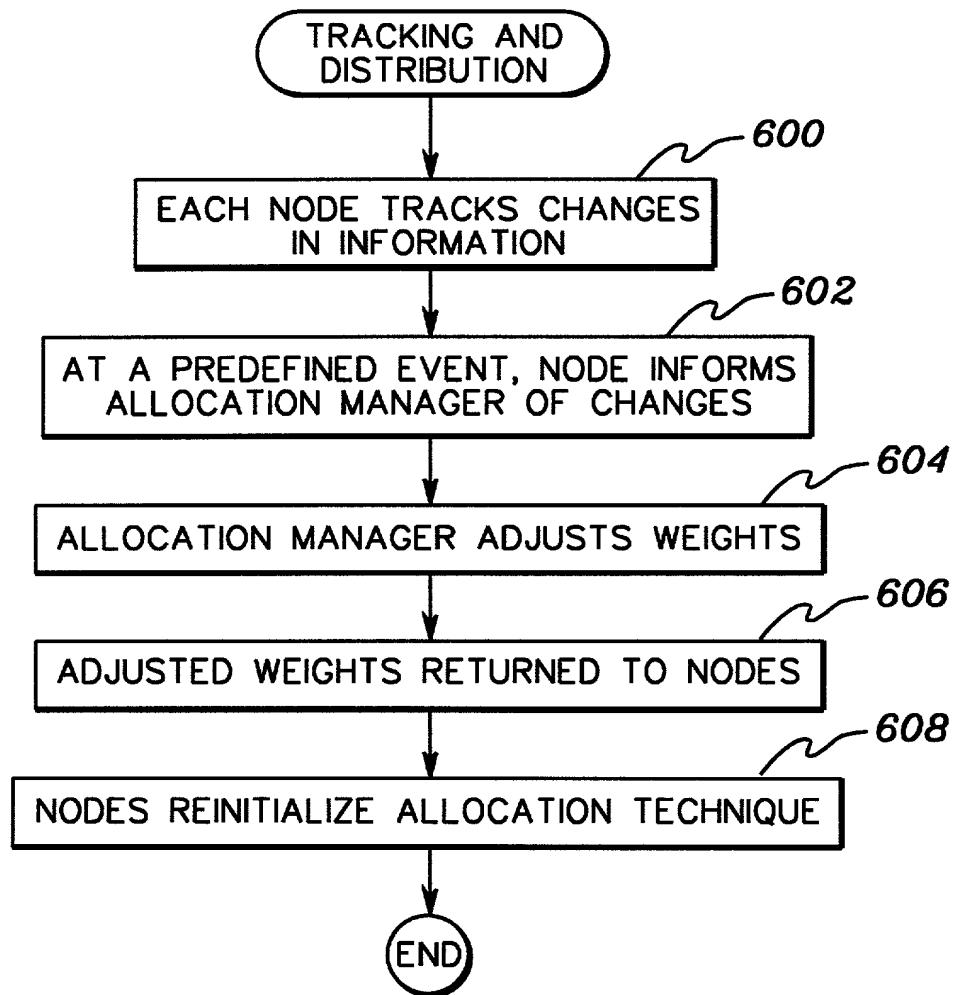
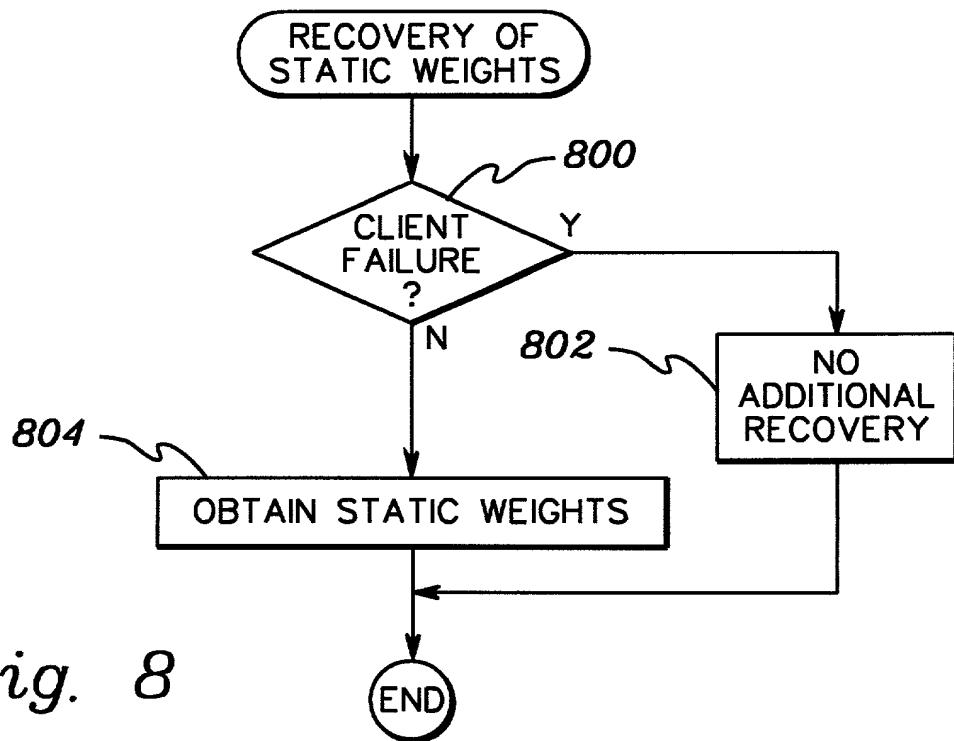
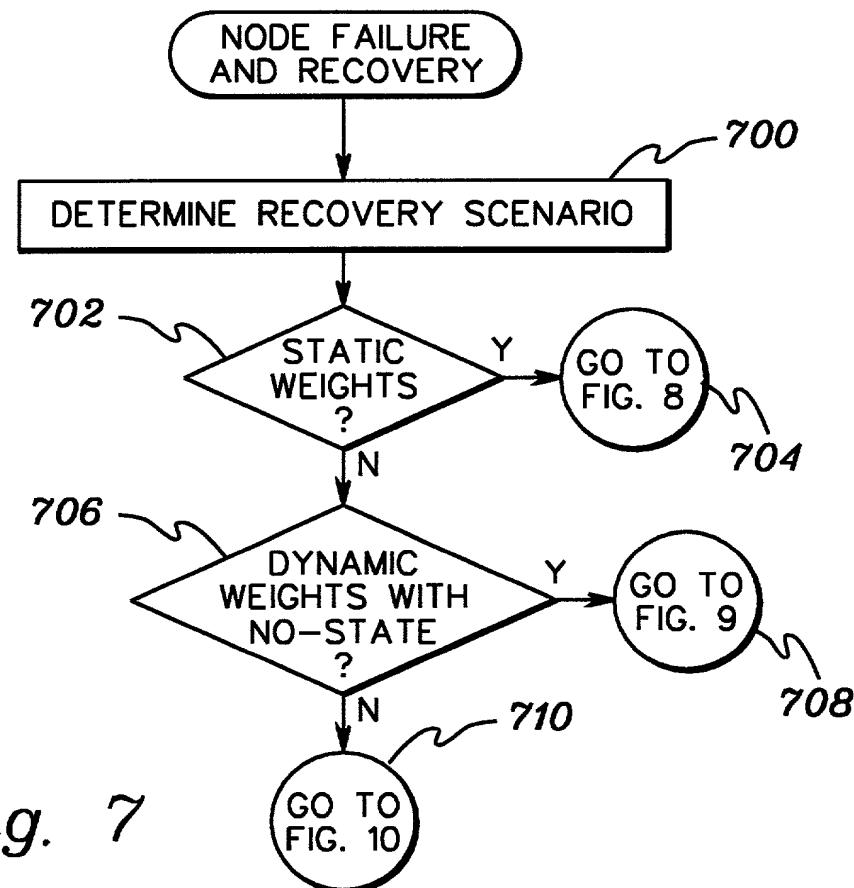


fig. 6



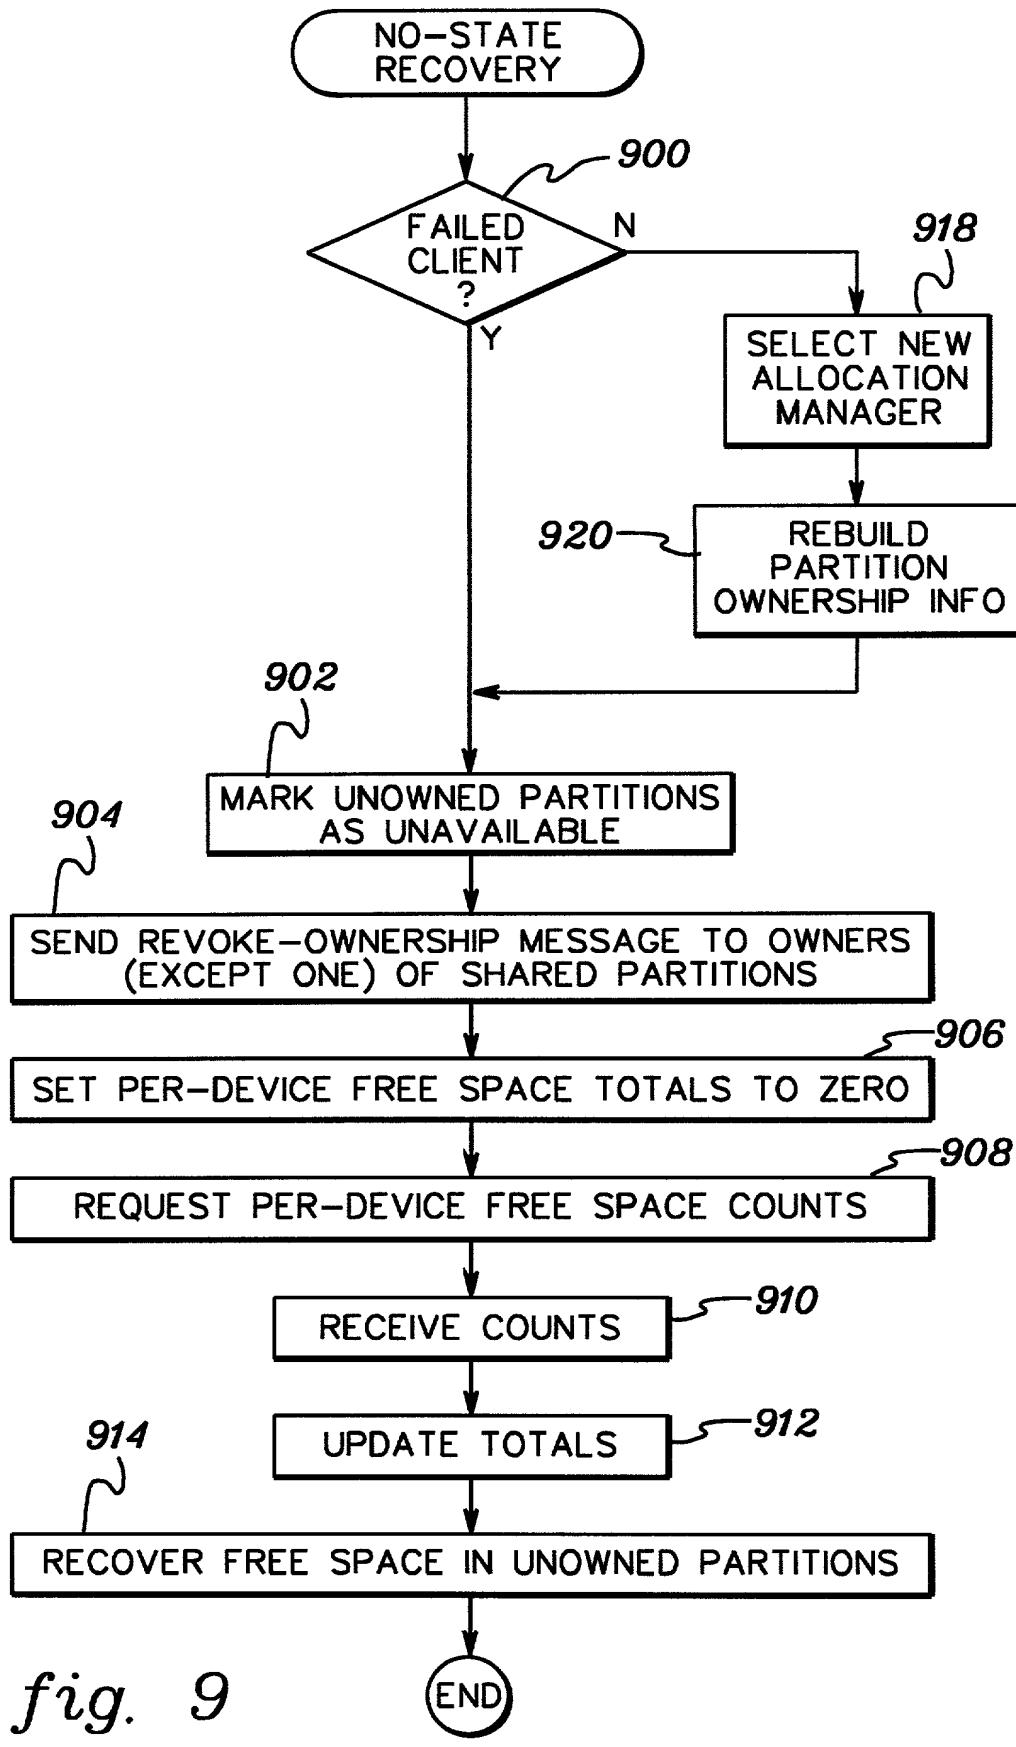


fig. 9

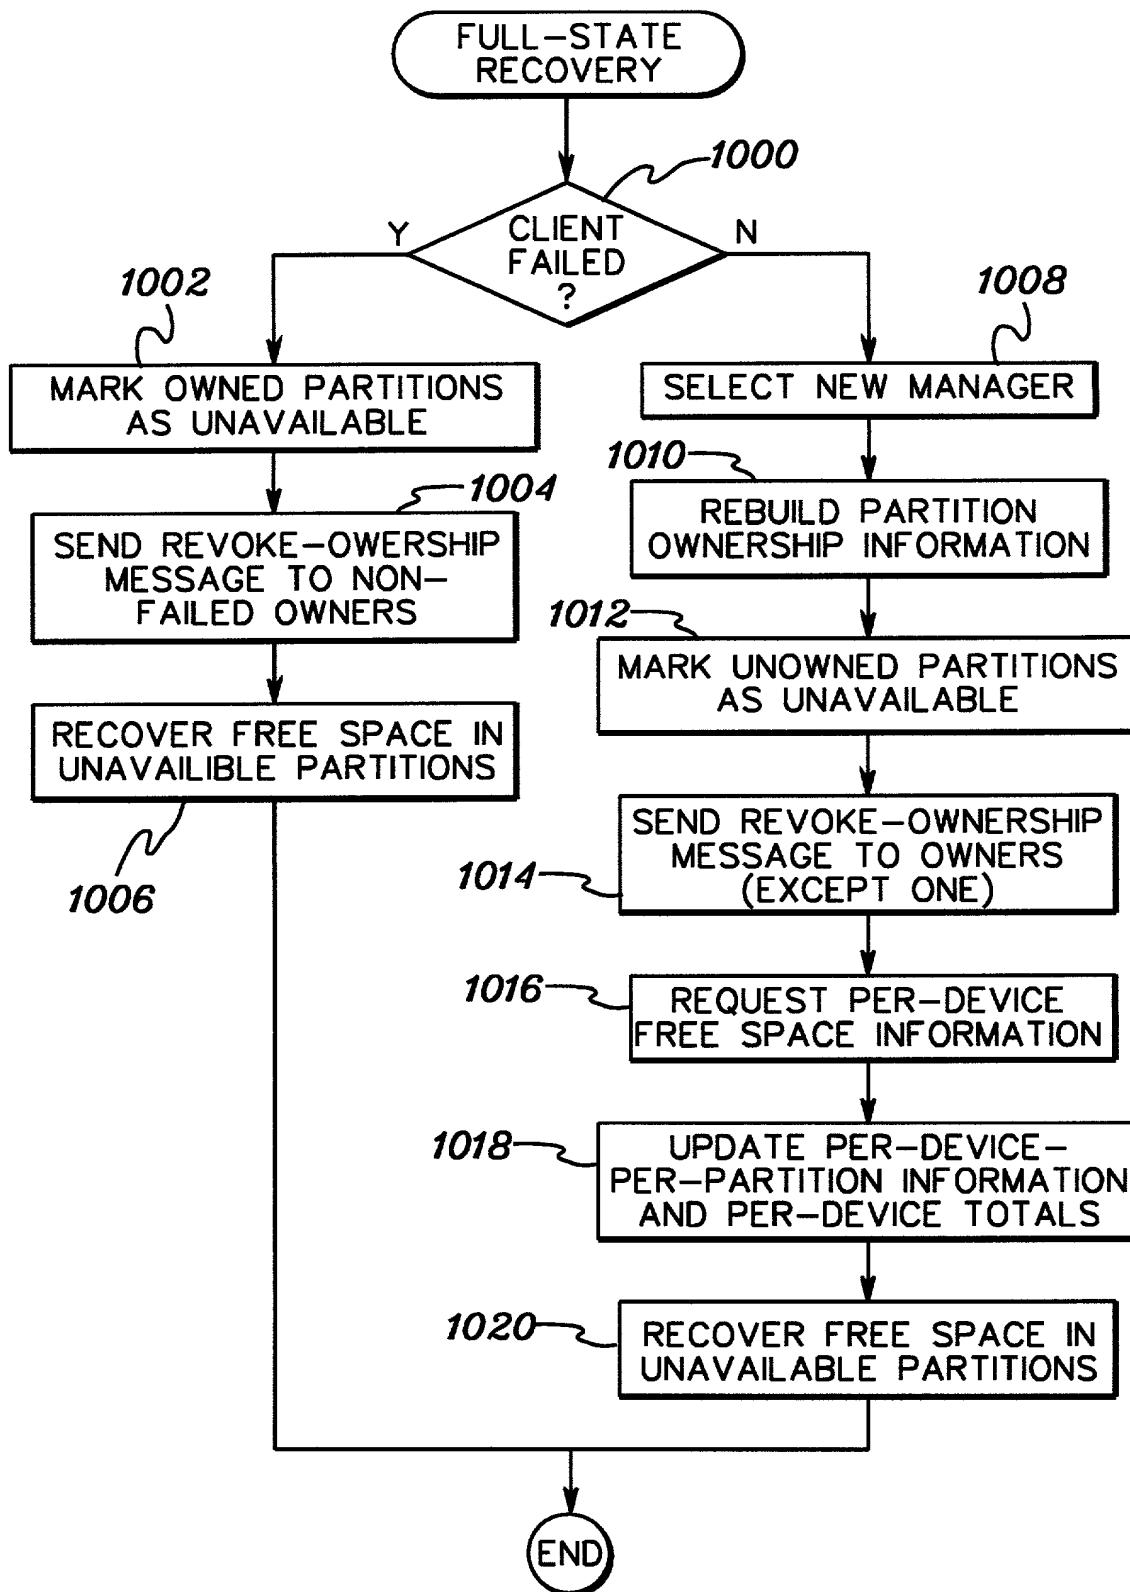


fig. 10

Docket No.
POU9-2000-0112-US1

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

A PLURALITY OF FILE SYSTEMS USING WEIGHTED ALLOCATION TO ALLOCATE SPACE ON ONE OR MORE STORAGE DEVICES

the specification of which

(check one)

is attached hereto.

was filed on _____ as United States Application No. or PCT International Application Number _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>
(Number)	(Country)	(Day/Month/Year Filed)	<input type="checkbox"/>

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional application(s) listed below:

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)

(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (*list name and registration number*)

Lynn L. Augspurger, Reg. No. 24,227

Christopher A. Hughes, Reg. No. 26,914

Lawrence D. Cutter, Reg. No. 28,501

Edward A. Pennington, Reg. No. 32,588

Marc A. Ehrlich, Reg. No. 39, 966

John E. Hoel, Reg. No. 26,279

William B. Porter, Reg. No. 33,135

Joseph C. Redmond, Jr., Reg. No. 18,753

Floyd A. Gonzalez, Reg. No. 26,732

Jeff Rothenberg, Reg. No. 26,429

William A. Kinnaman, Jr., Reg. No. 27,650

Kevin P. Radigan, Reg. No. 31,789

Lily Neff, Reg. No. 38,254

Blanche E. Schiller, Reg. No. 35,670

Andrew J. Woznicki, Jr., Reg. No. 43,995

Send Correspondence to: **Blanche E. Schiller, Esq.**
HESLIN & ROTHENBERG, P.C.
5 Columbia Circle
Albany, NY 12203

Direct Telephone Calls to: (*name and telephone number*)
Blanche E. Schiller, Esq. (518) 452-5600

Full name of sole or first inventor WAYNE A. SAWDON	Date
Sole or first inventor's signature	
Residence 6362 San Anselmo Way, San Jose, CA 95119	
Citizenship United States of America	
Post Office Address 6362 San Anselmo Way, San Jose, CA 95119	

Full name of second inventor, if any ROGER L. HASKIN	Date
Second inventor's signature	
Residence 13585 Struzenberg Court, Morgan Hill, CA 95037	
Citizenship United States of America	
Post Office Address 13585 Struzenberg Court, Morgan Hill, CA 95037	

Full name of third inventor, if any FRANK B. SCHMUCK	
Third inventor's signature	Date
Residence 406-A Union Avenue, Campbell, CA 95008	
Citizenship German	
Post Office Address 406-A Union Avenue, Campbell, CA 95008	

Full name of fourth inventor, if any JAMES C. WYLLIE	
Fourth inventor's signature	Date
Residence 18392 Chadbourne Lane, Monte Sereno, CA 95030	
Citizenship United States of America	
Post Office Address 18392 Chadbourne Lane, Monte Sereno, CA 95030	

Full name of fifth inventor, if any	
Fifth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	

Full name of sixth inventor, if any	
Sixth inventor's signature	Date
Residence	
Citizenship	
Post Office Address	